arXiv:1607.02804v2 [stat.ME] 19 Aug 2016

Estimating the number of species to attain sufficient representation in a random sample

Chao Deng and Andrew D. Smith

Molecular and Computational Biology
University of Southern California, Los Angeles, California

August 23, 2016

Abstract

The statistical problem of using an initial sample to estimate the number of unique species in a larger sample has found important applications in fields far removed from ecology. Here we address the general problem of estimating the number of species that will be represented at least r times, for any $r \geq 1$, in a future sample. We derive a procedure to construct estimators that apply universally for a given population: once constructed, they can be evaluated as a simple function of r. Our approach is based on a relationship between the number of species represented at least r times and the higher derivatives of the average number of unique species discovered per unit of sampling. We further show the estimators retain asymptotic behaviors that are essential for applications on large-scale data sets. We evaluate the performance of this approach by both simulation and applying it to analyze Dickens' vocabulary, the topology of a Twitter social network, and DNA sequencing data.

A random sample of N individuals is drawn from a population. Each individual belongs to exactly one species, and the total number L of species in the population is finite but not known. Let n_j be the number of species represented by exactly j individuals in this sample, so that

$$N = \sum_{j \geq 1} j n_j.$$

The number S_r of species represented r or more times in the initial sample is

$$S_r = \sum_{j \geq r} n_j.$$

Imagine that a second larger sample of size tN is drawn from the same population. The sampling effort $t > 1$ should bring to mind a "scaled up" experiment. This second sample may take the form of an expansion of the initial sample, but may also be a separate sampling experiment from the same underlying population. We are concerned with estimating the expected number $\mathrm{E}(S_r(t))$ of species represented at least r times in the second sample.

Related inference problems have been the focus of much statistical development, with canonical applications in linguistics and ecology (Fisher et al., 1943; Efron & Thisted, 1976; Zipf, 1935). The task of estimating L, often called "species richness," is well known and has found a broad diversity of applications (Harwit & Hildebrand, 1986; Stam, 1987; Kroes et al., 1999; Van Der Heijden et al., 2003; Mora et al., 2011; Deng et al., 2015). When plotted as a function of t, the quantity $S_1(t)$ is called the "accumulation curve"

(Colwell et al., 2004). As t approaches infinity the accumulation curve converges to L if every species is observable, and the same behavior clearly holds for each $S_r(t)$.

The quantity $S_1(t)$ may not be sufficient when the interests of the study are "common species" (Preston, 1948; Pearman & Weber, 2007). The parameter $r > 1$ in $S_r(t)$ is naturally applied to distinguish commonness from rarity. For example Jaccard (1912) considered common species to be those that were captured in at least four areas (species represented at least 4 times). Chazdon et al. (1998) defined abundant species as those represented by more than 10 individuals. The concept of "common species" also appears on other fields. Huberman et al. (2008) only concerned Twitter users with at least two posts, whom were called "active users." Google Scholar uses the number of publications cited at least 10 times by others (i10-index) to measure scholarly influence. Ng et al. (2010) filtered out single nucleotide polymorphisms (SNPs) covering fewer than 8 sequenced reads in a DNA sequencing experiment. In each of these cases, the value of r was established using application-specific criteria and in each $r > 1$.

The mixture of Poisson distributions has been widely used in capture-recapture experiments (Fisher et al., 1943; Good & Toulmin, 1956; Efron & Thisted, 1976). The Poisson distribution models the number of individuals for a species, with a latent distribution capturing varieties among species. Because the populations of interest are almost never homogeneous, accurate inference must somehow address the latent distribution of relative species abundance. As a notable early example, (Fisher et al., 1943) assumed species abundance followed a latent gamma distribution. Although other parametric distributions have been investigated (Bulmer, 1974; Sichel, 1975; Burrell & Fenton, 1993), there are problems with using parametric distributions in practice. In many cases no simple parametric form is suitable to explain the data. In other cases, distinct parametric forms may appear to fit the data equally well but exhibit very different extrapolation behaviors (Engen, 1978).

Good & Toulmin (1956) established a non-parametric empirical Bayes framework that served as the foundation for much subsequent non-parametric methodology (Efron & Thisted, 1976; Gale & Sampson, 1995; Daley & Smith, 2013). Assume the number of individuals from species i in a sample follows a Poisson process with rates λ_i per unit sampling effort, which follow an unknown latent distribution $G(\lambda)$. Good and Toulmin derived an estimator for the expected value of $S_1(t)$ while avoiding direct inference of $G(\lambda)$ (Good & Toulmin, 1956). This estimator takes the form of a power series with coefficients based on the count frequency n_j from the initial sample. Unfortunately, the Good-Toulmin estimator diverges for $t > 2$, and is therefore almost useless in most modern applications.

Daley & Smith (2013) solved the divergence problem by applying rational function approximation (RFA) to the Good-Toulmin power series. This approach was motivated by applications associated with DNA sequencing libraries (Daley & Smith, 2013; Daley & Smith, 2014), where data sizes are many orders of magnitude larger than traditional applications. However, the approach of Daley & Smith (2013) does not directly extend to $r > 2$ because of the intricacy of the general expression of the Good-Toulmin power series (Daley, 2014).

Here we describe a new approach to estimate the number of species that will be represented at least r times in a sample of size tN, based on a sample of size N from the same population. This approach comes from two insights. The first is a relationship between the values we seek to estimate and the higher-order derivatives of the discovery rate, defined as the expected number $\mathrm{E}(S_1(t))$ of species divided by the sampling effort t. The latter quantity is the average expected number of species per unit of sampling effort. The second observation is a convenient numerical algorithm suggested by this relationship. This algorithm generates estimators that, once constructed, can be directly evaluated as a function of r. We show that these estimators satisfy essential limiting properties as either r or t approaches infinity, providing a theoretical basis for the accuracy and stability we demonstrate in large-scale applications to linguistic, social network

and DNA sequencing data.

1 Sampling model and background

As noted above, we use n_j to denote the number of species observed j times for one unit of sampling effort, $j = 1, 2, \ldots$. Clearly n_0 is unobserved. Let $n_j(t)$ be the random variable whose value is the number of species observed exactly j times for sampling effort t. The number $S_r(t)$ of species observed at least r times as a function of t can then be defined as

$$S_r(t) = \sum_{j=r}^{\infty} n_j(t) = S_1(t) - \sum_{j=1}^{r-1} n_j(t). \tag{1}$$

Our goal is to estimate $\mathrm{E}(S_r(t))$ using information from the n_j. Here we review results from the non-parametric empirical Bayes framework of Good & Toulmin (1956); Efron & Thisted (1976) that are needed for our derivation.

Recall that $G(\lambda)$ is the cumulative distribution function associated with the Poisson rates λ_i, for $i = 1, 2, \ldots, L$. The expected value of $n_j(t)$ is

$$\mathrm{E}(n_j(t)) \;=\; L \int \frac{(\lambda t)^j e^{-\lambda t}}{j!} dG(\lambda) \;=\; t^j L \int \frac{\lambda^j}{j!} e^{-\lambda} e^{-\lambda(t-1)} dG(\lambda).$$

Replacing $\exp(-\lambda(t-1))$ with its power series yields

$$\begin{aligned}
\mathrm{E}(n_j(t)) &= t^j L \int \frac{\lambda^j}{j!} e^{-\lambda} \sum_{i=0}^{\infty} \frac{(-\lambda(t-1))^i}{i!} dG(\lambda) \\
&= t^j \sum_{i=0}^{\infty} (-1)^i \binom{i+j}{j} (t-1)^i L \int \frac{\lambda^{i+j}}{(i+j)!} e^{-\lambda} dG(\lambda).
\end{aligned}$$

Since the expected number of species observed j times in the initial sample is

$$\mathrm{E}(n_j) = L \int \frac{\lambda^j e^{-\lambda}}{j!} dG(\lambda),$$

the expected value of $n_j(t)$ can be expressed as

$$\mathrm{E}(n_j(t)) = t^j \sum_{i=0}^{\infty} (-1)^i \binom{i+j}{j} (t-1)^i \mathrm{E}(n_{i+j}), \tag{2}$$

which is due to Good & Toulmin (1956). Similarly, the expected number of unique species for sampling effort t, known as Good-Toulmin power series (Good & Toulmin, 1956, eqn. 19), is

$$\mathrm{E}(S_1(t)) = \sum_{j=1}^{\infty} \mathrm{E}(n_j) + \sum_{j=1}^{\infty} (-1)^{j+1} (t-1)^j \mathrm{E}(n_j). \tag{3}$$

Substituting n_j for $\mathrm{E}(n_j)$ above yields an estimator, but one that diverges for $t > 2$.

2 Number of species with sufficient representation

The derivative of $\mathrm{E}(S_1(t))$, defined in Equation (3), with respect to t is

$$\frac{d}{dt}\mathrm{E}(S_1(t)) = \sum_{j=1}^{\infty}(-1)^{j+1}j(t-1)^{j-1}\mathrm{E}(n_j) = \sum_{j=0}^{\infty}(-1)^{j}(j+1)(t-1)^{j}\mathrm{E}(n_{j+1}).$$

To generalize, the j^{th} derivative of $\mathrm{E}(S_1(t))$ is

$$\frac{d^j}{dt^j}\mathrm{E}(S_1(t)) = (-1)^{j-1}j!\sum_{i=0}^{\infty}(-1)^{i}\binom{i+j}{j}(t-1)^{i}\mathrm{E}(n_{i+j}).$$

By comparing the above expression with the power series for $\mathrm{E}(n_j(t))$ in Equation (2), we obtain

$$\mathrm{E}(n_j(t)) = \frac{(-1)^{j-1}t^j}{j!}\,\frac{d^j}{dt^j}\mathrm{E}(S_1(t)), \tag{4}$$

which has been noted previously (Kalinin, 1965). The above derivation can only be applied when $t < 2$, but one can verify that the relation also holds for $t \geq 2$. Equation (4) shows how to replace the $\mathrm{E}(n_j(t))$ in Equation (1) with the j^{th} derivative of $\mathrm{E}(S_1(t))$. The following relationship is the foundation for our estimator; a proof is given in the appendix.

Proposition 1. *Let $S_r(t)$ denote the number of species observed at least r times in a sample of size tN. For any positive integer r,*

$$\mathrm{E}(S_r(t)) = \frac{(-1)^{r-1}t^r}{(r-1)!}\,\frac{d^{r-1}}{dt^{r-1}}\left(\frac{\mathrm{E}(S_1(t))}{t}\right). \tag{5}$$

The quantity $\mathrm{E}(S_1(t))/t$ in Equation (5) contains information sufficient for estimating $\mathrm{E}(S_r(t))$. We call this ratio the discovery rate as $E(S_1(t))/t$ reflects the average efficiency of obtaining new species per unit of sampling effort. The discovery rate can be used as one characterization for diversity in a population. One extreme case is seen when the discovery rate is a constant. By Equation (5) this implies $\mathrm{E}(S_r(t))$ is equal to 0 for $r > 1$. The statement is true when $\mathrm{E}(S_1(t)) = tN$, which means every individual is expected to be a new species, with the discovery rate exactly equal to 1.

3 A universal estimator and estimation algorithm

Proposition 1 tells us that any estimator for $\mathrm{E}(S_1(t))/t$ that is defined over $(1, \infty)$ and is of differentiability class C^{r-1} gives an estimator for $\mathrm{E}(S_r(t))$. Good & Toulmin (1956) derived a widely known estimator for $\mathrm{E}(S_1(t))$ as alternating power series:

$$\mathrm{E}(S_1(t)) \;=\; \sum_{j=1}^{\infty}\mathrm{E}(n_j) + \sum_{j=1}^{\infty}(-1)^{j-1}(t-1)^{j}\mathrm{E}(n_j) \;\approx\; \sum_{j=1}^{\infty}n_j + \sum_{j=1}^{\infty}(-1)^{j-1}(t-1)^{j}n_j,$$

as the observed n_j is an unbiased estimate of $\mathrm{E}(n_j)$. We already discussed the problems with this power series when $t > 2$. In this section we show how an RFA representation for the discovery rate takes a particular form that dramatically simplifies the estimator suggested by Proposition 1.

First we need a power series for the discovery rate centered at $t = 1$. The expansion of $1/t$, at $t = 1$, is

$$\frac{1}{t} = \frac{1}{1+t-1} = 1 - (t-1) + (t-1)^2 - \cdots$$

for $0 < t < 2$. Replacing $\mathrm{E}(S_1(t))$ with its power series yields

$$\begin{aligned}\mathrm{E}(S_1(t))/t &= \left(\sum_{j=1}^{\infty} \mathrm{E}(n_j) + \mathrm{E}(n_1)(t-1) - \mathrm{E}(n_2)(t-1)^2 + \cdots\right)\left(1 - (t-1) + (t-1)^2 + \cdots\right) \\ &= \sum_{j=1}^{\infty} \mathrm{E}(n_j) - (t-1)\sum_{j=2}^{\infty} \mathrm{E}(n_j) + (t-1)^2 \sum_{j=3}^{\infty} \mathrm{E}(n_j) - (t-1)^3 \sum_{j=4}^{\infty} \mathrm{E}(n_j) + \cdots\end{aligned}$$

The observed n_j are unbiased estimates of $\mathrm{E}(n_j)$, so the discovery rate can be estimated as

$$\hat{\mathrm{E}}(S_1(t))/t = \sum_{j=1}^{\infty} n_j - (t-1)\sum_{j=2}^{\infty} n_j + (t-1)^2 \sum_{j=3}^{\infty} n_j - (t-1)^3 \sum_{j=4}^{\infty} n_j + \cdots \tag{6}$$

Next we apply RFA to Equation (6). We restrict our attention to rational functions with numerator $A(t)$ of order $M-1$ and denominator $B(t)$ of order M:

$$\hat{\mathrm{E}}(S_1(t))/t \approx \frac{A(t)}{B(t)} = \frac{a_0 + a_1 t + a_2 t^2 + \cdots + a_{M-1} t^{M-1}}{b_0 + b_1 t + b_2 t^2 + \cdots + b_M t^M}. \tag{7}$$

We assume all RFAs are Padé approximants, which have the property that their Taylor expansion is equal to the power series up to order $2M-1$ (Baker Jr, 1961). Substituting $\hat{\mathrm{E}}(S_1(t))/t$ into Equation (5), our estimator is

$$\hat{\mathrm{E}}(S_r(t)) = \frac{(-1)^{r-1} t^r}{(r-1)!} \frac{d^{r-1}}{dt^{r-1}} \left(\frac{A(t)}{B(t)}\right). \tag{8}$$

Let $x_1, x_2, \ldots, x_M$ be the (complex) roots of the denominator $B(t)$ in Equation (7). If all x_i are distinct and none are equal to 0, we can express $\hat{\mathrm{E}}(S_1(t))/t$ as

$$\hat{\mathrm{E}}(S_1(t))/t = \frac{c_1}{t - x_1} + \frac{c_2}{t - x_2} + \cdots + \frac{c_M}{t - x_M}, \tag{9}$$

where the c_i, for $1 \le i \le M$, are coefficients of a partial fraction decomposition for (7). Note that because of our requirements on the degrees of $A(t)$ and $B(t)$ the partial fraction lacks a constant term. By substituting this expansion of Equation (9) for $A(t)/B(t)$ in Equation (8),

$$\hat{\mathrm{E}}(S_r(t)) = \frac{(-1)^{r-1} t^r}{(r-1)!} \frac{d^{r-1}}{dt^{r-1}} \left(\frac{c_1}{t - x_1} + \frac{c_2}{t - x_2} + \cdots + \frac{c_M}{t - x_M}\right).$$

The above expression simplifies by noticing that the derivatives take a convenient form:

$$\frac{d^{r-1}}{dt^{r-1}} \left(\frac{c_i}{t - x_i}\right) = (-1)^{r-1}(r-1)! \left(\frac{c_i}{(t - x_i)^r}\right).$$

We have arrived at our estimator for the number of species observed at least r times in sampling effort t:

$$S_r^{\star}(t) = \hat{\mathrm{E}}(S_r(t)) = \sum_{i=1}^{M} c_i \left(\frac{t}{t - x_i}\right)^r. \tag{10}$$

To summarize, the steps in constructing the estimators are:

1. Build the power series for $\hat{\mathrm{E}}(S_1(t))/t$ given by (6) using the count frequencies n_j from the initial sample of size N.

2. Compute an RFA given by (7) to the power series $\hat{\mathrm{E}}(S_1(t))/t$, estimating coefficients a_i for $0 \leq i < M$ and b_i for $0 \leq i \leq M$.

3. Decompose the RFA as partial fractions (9), estimating roots x_i and coefficients c_i, for $1 \leq i \leq M$.

4 Conditions for estimator stability in practice

Large-scale applications will require estimators that remain accurate over long-range extrapolation, and that exhibit stability as a function of r. The numerical computations in constructing our estimator make it difficult to guarantee all desired behaviors. However, two asymptotic properties of $\mathrm{E}(S_r(t))$ are especially important to retain in estimators:

(i) For any $r \geq 1$, the expected number $\mathrm{E}(S_r(t))$ of species observed at least r times converges to a positive constant as t approaches infinity.

(ii) For any $t \geq 1$, the expected number $\mathrm{E}(S_r(t))$ of species observed at least r times converges to 0 as r approaches infinity.

In practice the estimator $S_r^\star(t)$, constructed using the procedure of the previous section, may violate properties *(i)* and *(ii)*. Below we explain sufficient conditions for these essential properties to hold in the estimator, and how we use these sufficient conditions.

The constant limit indicated in property *(i)* is clearly the number of species L in the population. The estimator $S_r^\star(t)$ should also converge with t to a constant limit independent of r or the estimates will be inconsistent. A similar constraint for the special case $r = 1$ is mentioned by Boneh et al. (1998). From Equation (10), as t approaches infinity the limit of $S_r^\star(t)$ is the sum of all coefficients c_i, which we denote by $\hat{L}$. By the optimality property of Padé approximants, our estimate $\hat{\mathrm{E}}(S_1(t))$ is equal to S_1 at $t = 1$. If $S_r^\star(t)$ is non-decreasing, the constant limit $\hat{L}$ will be positive, and the estimator will satisfy property *(i)*.

Now we propose a sufficient condition for asymptotic property *(ii)*. Note that $Re(x)$ denotes the real part of complex number x.

Proposition 2. *Suppose M is a positive integer. Let c_i and x_i be complex numbers, and define*

$$S_r^\star(t) = \sum_{i=1}^{M} c_i \left(\frac{t}{t - x_i} \right)^r,$$

where r is a positive integer and $t \geq 1$. If $x_i \neq 0$ and $Re(x_i) \leq 0$, for $1 \leq i \leq M$, then $S_r^\star(t)$ converges to 0 as r goes to infinity, for any $t \in [1, \infty)$.

Proof. If $Re(x_i) \leq 0$ then for any given $t \in (1, \infty)$ and any x_i,

$$\|t - x_i\|^2 = t^2 - (x_i + \bar{x}_i)t + x_i \bar{x}_i \geq t^2.$$

Because we assume $x_i \neq 0$, the above inequality is strict. Therefore $\|t/(t - x_i)\| < 1$, and

$$\lim_{r \to \infty} \left\| \sum_{i=1}^{M} c_i \left(\frac{t}{t - x_i} \right)^r \right\| \;\leq\; \sum_{i=1}^{M} \lim_{r \to \infty} \|c_i\| \left\| \frac{t}{t - x_i} \right\|^r \;=\; 0. \qquad \square$$

The selection of the free parameter M, which controls the degree of the RFA, can be used to identify estimators that exhibit the asymptotic properties *(i)* and *(ii)*. For each M and corresponding unique optimal RFA $A(t)/B(t)$, we test: (1) whether the resulting estimator $\hat{\mathrm{E}}(S_1(t))$, noting that here $r = 1$, is an increasing function; and (2) whether the roots of x_i satisfy the antecedents of Proposition 2. Then among values of M for which these two tests pass, we choose the maximum. In theory M should be selected sufficiently large so that the RFA captures the underlying complexity of the latent distribution $G(\lambda)$. The danger in selecting very large values of M is that the procedure will use observed n_j for larger j. In practice those will show greater deviation from the $\mathrm{E}(n_j)$ that they estimate.

5 Simulation

We use simulation to evaluate the performance of the estimator $S_r^\star(t)$. According to our hypothesis, the estimator $S_r^\star(t)$ should work appropriately for $r > 1$ if it can accurately estimate the discovery rate or the species accumulation curve (the case $r = 1$). To test this hypothesis, we compared the expected number of species represented at least r times with the estimates for multiple r under three heterogeneous population models: negative binomial, Zipf and Zipf-Mandelbrot.

Our simulation scheme is inspired by Chao & Shen (2004). The total number of species in the simulation is 100,000. The expected size of initial samples is around 180,000. In each population, the number of individuals captured for species i in a sample follows a Poisson distribution with rates λ_i, for $1 \leq i \leq 100000$. The rates λ_i are generated by different distributions to simulate populations with different heterogeneity levels.

- Negative binomial: the λ_i are generated from a gamma density

$$f(x) = \frac{1}{\Gamma(\alpha)\beta^\alpha} x^{\alpha-1} e^{-\frac{x}{\beta}}$$

  with $\alpha = 0.5$, $\beta = 3.6$ (CV $= 1.41$, $\mathrm{E}(n) = 180000$).

- Zipf: $\lambda_i = 20000/(i + 10)$ for $i = 1, 2, \ldots, 100000$ (CV $= 10.60$, $\mathrm{E}(n) = 183226$).

- Zipf-Mandelbrot: $\lambda_i = 150000/(y + 20)^{1.31}$ for $i = 1$ to 100000 (CV $= 17.94$, $\mathrm{E}(n) = 186225$).

For each model we simulated 1000 datasets as initial samples, which were used to construct the estimators $S_r^\star(t)$. Here we only consider $t \in (1, 20]$ since the range is mostly sufficient in practical applications. We measure the overall performance by the scaled error, defined as the L_2 norm between the expected and estimated accumulation curves, divided by the L_2 norm of the expected curve. This quantity is approximated by 19 discrete points along the curve. For the case $r = 1$, the estimator $S_r^\star(t)$ accurately predicted the species accumulation curves on all three models. The scaled errors were 7.02% for the negative binomial model, 7.16% for the Zipf model, and 2.89% for the Zipf-Mandelbrot model. The coefficient of variation (CV) with respect to $(\lambda_1, \lambda_2, \ldots, \lambda_S)$ is defined as

$$\frac{[S^{-1} \sum_{i=1}^{S} (\lambda_i - \bar{\lambda})]^{1/2}}{\bar{\lambda}},$$

where $\bar{\lambda} = S^{-1} \sum_{i=1}^{S} \lambda_i$. A large CV indicates a high level of heterogeneity in the population.

The estimator $S_r^\star(t)$ maintained high accuracy for both the Zipf and the Zipf-Mandelbrot assumptions when $r < 100$ (Figure 1). Even for r as large as 99, the scaled errors were only 12.08% and 16.8%

respectively. However, the estimator $S_r^\star(t)$ did not retain the same accuracy on the negative binomial model. The scaled error was much larger compared with the other two models. For example at $r = 20$, the scaled error was less than 5% on both the Zipf model and Zipf-Mandelbrot model while it was 32.60% on the negative binomial model (Figure 1). The result suggests the estimator $S_r^\star(t)$ favors the highly heterogeneous population, which in turn has a large CV.

The standard errors are reasonable when $1 < r <\leq 20$. They are less than 10% with respect to their expected values. For a fixed r, we found the standard error increased as a function of t (Tables 1–3). It is intuitive since estimates should be less certain when extrapolating further. We also observed that in general the estimates were conservative (Tables 1–3). For those cases where the estimates were greater than the expected values, the difference is less than 5% of the expected values. For instance, when $t = 20$ and $r = 2$, the expected value and the estimate were 98167 and 99514 in the Zipf model (Table 2). The estimate overestimated the true value by 1347, which is around 1.4% of the true value. Being accurate and conservative are very useful when estimating sample size needed to observe some portion of the population diversity in an experiment.

Practical non-parametric methods for estimating $S_r(t)$ when $r > 1$ are virtually non-existent in the literature. We therefore selected two parametric approaches for comparison: the zero-truncated negative binomial and zero-truncated Poisson respectively, and a nonparametric approach suggested by Chao & Shen (2004) and compared them for the case $r = 1$. The Poisson distribution is commonly used to model the occurrences of an events. The negative binomial distribution is a natural alternative to the Poisson when counts are overdispersed, and the zero-truncated version has a long history in modeling species capture data (Fisher et al., 1943). The result is in the supplementary (Section S5).

6 Applications

Here we present results from applying the estimator to data from three different application domains. In each application, we subsample from the entire dataset to simulate an initial experiment. Then we apply the proposed approach to estimate $\mathrm{E}(S_r(t))$. We compare the estimates with the expected values generated by subsampling the entire dataset.

We do not have an expression for the variance of $S_r^\star(t)$ that can account for the effects from each step in our procedure. We approximate the variance of our estimates by bootstrapping the initial sample (details in Supplementary Information). Often a normal distribution is assumed for the purpose of generating confidence intervals. However, Daley & Smith (2013) reported a skew to the confidence interval upward for the rational function approximation for the case $r = 1$. Therefore we follow their procedure and construct the confidence interval through the log-normal approach (Chao, 1987). Additional information for confidence interval is in the supplementary (Section S7)

6.1 Linguistics

We applied the estimator $S_r^\star(t)$ to infer word frequencies in a sample of Charles Dickens' works. We used data from Project Gutenberg (Hart, 1971) that is included in the R package *zipfR* (Evert & Baroni, 2007). This data set contains roughly 2.8M written words, of which just over 41k are distinct. We sampled 300k words from the dataset as an initial sample and applied $S_r^\star(t)$ for values of r between 1 and 20. Figure 2 shows the estimated curves along with actual curves from the entire Dickens data set. The estimated curves track the corresponding actual curves very closely. Table 4 shows estimated and actual values for extrapolations of $5\times$ and $9\times$, the latter being the maximum possible given the size of the data set. Even

at $r = 20$ the scaled error never exceeds 10% We also examined the behavior of $S_r^\star(t)$ as a function of r. As can be seen from Figure 2(b), $S_r^\star(t)$ remains accurate for large values of r, starting to diverge only after $r \approx 70$.

6.2 A social network

We also applied the estimator $S_r^\star(t)$ to predict the number of Twitter users that will have r or more followers, based on a small sample of "(follower, followed)" relations. We obtained a data set from the Social Computing Data Repository (Zafarani & Liu, 2009). This data set contains 11.3M users and over 85.3M following relationships, which form edges in this social network. We sampled 5M edges as an initial sample and used this to estimate the number of users with at least r followers in larger sets of following relationships. Estimates using $S_r^\star(t)$ showed high accuracy when extrapolating to $5\times$ the size of the initial sample, as can be seen in Figure 3(a). For example, when the total number of following relations is 25M, we should see roughly 1.2M individuals with at least 2 followers. Our prediction of just over 1.3M was off by 6%. Accuracy decreases for larger extrapolations. The entire dataset contains 2.8M users that each have at least $r = 2$ followers, but our estimator predicted 3.5M, an overestimate of 23%. Interestingly, accuracy does not rapidly worsen with r (Table 5) and seems to remain high for values of r up to 100 (Figure 3(b)), consistent with our results on the linguistic data set.

We used the zero truncated negative binomial (ZTNB) to estimate $\mathrm{E}(S_r(t))$ for both the Dickens and Twitter data sets. The performance is measured by the scaled error, which is the absolute value of subtraction of the observed value and the estimate divided by the observed value. In both cases, the error from ZTNB was substantially higher than from our estimator (Figure 4). The estimates from ZTNB are less sensitive to r comparing with $S_r^\star(t)$, which causes the large scaled error (Figure 5).

6.3 DNA sequencing

To evaluate our approach on an larger scale, we applied it to predict the number of base pairs in the human genome that will be represented at least r times in sequencing data set. In genomics terminology, these are the positions in the genome covered by at least r sequenced reads, or the positions with coverage depth at least r. Coverage depth is critical in genetics studies, for example in detecting SNPs, where candidate SNPs with low coverage depth are often discarded. Knowing the distribution of coverage depth can help researchers in experimental design, informing the total amount of required sequencing in order to attain sufficient coverage over an acceptable number of genomic sites.

We used four publicly available DNA sequencing experiments to evaluate the performance of $S_r^\star(t)$ (Zong et al., 2012; Lu et al., 2012; Wang et al., 2012). For each dataset, we subsampled 5M reads as a "shallow" sequencing experiment, and we used this information to estimate the number of sites that would attain minimum required depth as sequencing continues (details in supplementary information). Figure 6(a) shows the curves for estimated values $S_r^\star(t)$ for multiple values of r, along with the actual expected values obtained by repeated sub-sampling from the full data set. Figure 6(b) presents the same information for estimates based on the ZTNB. This data set is sufficiently large to reveal the inflection points in the curves when $r > 1$. Estimates from $S_r^\star(t)$ closely track the true values. The ZTNB, on the other hand, overestimates $\mathrm{E}(S_1(t))$ and then under-estimates for other values of r. On the remaining three data sets, both methods show varying accuracy, which is almost always higher for $S_r^\star(t)$ (Figure S2(a)-(e)).

7 Discussion

We presented an approach to estimate the number of species that will be represented at least r times in a sample by using information in a small sample from the same population. The non-parametric estimators obtained by our approach are universal in that they apply across values of r for a given population. We have shown these estimators to have favorable properties in theory, and also to give highly accurate estimates in practice. Accuracy remains high for large values of r and for long range extrapolations. This approach builds upon the theoretical non-parametric empirical Bayes foundation of Good & Toulmin (1956), providing a practical way to compute estimates that are both accurate and stable.

The foundation for our approach is a relationship between the expected number $\mathrm{E}(S_r(t))$ of species represented r times and the $(r-1)^{\mathrm{th}}$ derivative of the discovery rate. This relationship characterizes $\mathrm{E}(S_r(t))$ directly, avoiding the summation of $\mathrm{E}(n_j(t))$ estimates. Clearly any estimator for either $S_r(t)$ or $n_r(t)$ provides a means of estimating both quantities. Equation (1) defines $S_r(t)$ as a sum of $n_j(t)$ for $j \geq r$. Similarly $n_r(t)$ can be written as $S_r(t) - S_{r+1}(t)$. We prefer to work with $S_r(t)$ in practice for the additional reason that $\mathrm{E}(S_r(t))$ is an increasing function – an extremely valuable property when diagnosing numerical issues during the construction.

Our empirical accuracy evaluations were based on applications where the underlying data sets can be considered high compared to traditional applications from ecology. While having more data naturally leads to better estimates, for our approach even to be applicable the data must meet certain requirements. The number of non-zero values for observed n_j from the initial sample dictates the maximum value for M, the degree of the RFA. Our implementation requires that the first four counts (n_1, n_2, n_3 and n_4) be non-zero. When more counts are available, and when they are sufficiently large, using a larger value for M will lead to better estimates. However, if for larger j the observed counts n_j are too small, those counts may not accurately reflect the moments of the underlying latent distribution. This dependence on having sufficient total data from the initial sample may limit the applicability of our estimator in traditional species abundance applications from ecology. However, for modern data analysis applications the initial samples are frequently in the millions. In those applications, if any of the first four counts are equal to zero, the initial sample is likely too small and any method would produce inaccurate results. Indeed for many applications that involve big data, two terms are rarely sufficient to account for heterogeneity in the underlying populations.

From the derivation, the estimator $S_r^\star(t)$ can be applied for any $t \geq 1$ and $r \geq 1$. We further extend the domain of t in the supplementary so that $S_r^\star(t)$ can be applied for $t > 0$. However, we observed that the standard error of $S_r^\star(t)$ increased as t increased, and in some cases the estimates become inaccuracy for large r. For general purposes we recommend users to apply $S_r^\star(t)$ to extrapolating no more than 20 times under $r \leq 20$, which help to control the accuracy and the standard error. From our experience, extrapolating to 20 times should be sufficient for most applications and estimates under this range usually satisfy the required precision.

References

Baker Jr GA (1961) Application of the Padé approximant method to the investigation of some magnetic properties of the Ising model. *Physical Review* 124:768.

Boneh S, Boneh A, Caron RJ (1998) Estimating the prediction function and the number of unseen species in sampling with replacement. *Journal of the American Statistical Association* 93:372–379.

Bulmer M (1974) On fitting the Poisson lognormal distribution to species-abundance data. *Biometrics* 30:101–110.

Burrell QL, Fenton MR (1993) Yes, the GIGP really does work – and is workable! *Journal of the American Society for Information Science* 44:61–69.

Chao A (1987) Estimating the population size for capture-recapture data with unequal catchability. *Biometrics* 43:783–791.

Chao A, Shen TJ (2004) Nonparametric prediction in species sampling. *Journal of agricultural, biological, and environmental statistics* 9:253–269.

Chazdon RL, Colwell RK, Denslow JS, Guariguata MR (1998) Statistical methods for estimating species richness of woody regeneration in primary and secondary rain forests of northeastern costa rica (ed. f. dallmeier & j. a. comiskey). *Paris: Parthenon Publishing* pp. 285–309.

Colwell RK, Mao CX, Chang J (2004) Interpolating, extrapolating, and comparing incidence-based species accumulation curves. *Ecology* 85:2717–2727.

Daley T (2014) Non-parametric models for large capture-recapture experiments with applications to DNA sequencing Ph.D. diss., University of Southern California.

Daley T, Smith AD (2013) Predicting the molecular complexity of sequencing libraries. *Nature Methods* 10:325–327.

Daley T, Smith AD (2014) Modeling genome coverage in single-cell sequencing. *Bioinformatics* 30:3159–3165.

Deng C, Daley T, Smith A (2015) Applications of species accumulation curves in large-scale biological data analysis. *Quantitative Biology* 3:135–144.

Efron B, Thisted R (1976) Estimating the number of unseen species: How many words did Shakespeare know? *Biometrika* 63:435–447.

Engen S (1978) *Stochastic Abundance Models* Chapman and Hall.

Evert S, Baroni M (2007) *zipfR*: Word frequency distributions in R In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, Posters and Demonstrations Sessions*, pp. 29–32, Prague, Czech Republic (R package version 0.6-6 of 2012-04-03).

Fisher RA, Corbet AS, Williams CB (1943) The relation between the number of species and the number of individuals in a random sample of an animal population. *The Journal of Animal Ecology* pp. 42–58.

Gale WA, Sampson G (1995) Good-Turing frequency estimation without tears. *Journal of Quantitative Linguistics* 2:217–237.

Good I, Toulmin G (1956) The number of new species, and the increase in population coverage, when a sample is increased. *Biometrika* 43:45–63.

Hart M (1971) *Project Gutenberg* Project Gutenberg.

Harwit M, Hildebrand R (1986) How many more discoveries in the universe? *Nature* 320:724–726.

Huberman B, Romero DM, Wu F (2008) Social networks that matter: Twitter under the microscope. *First Monday* 14.

Jaccard P (1912) The distribution of the flora in the alpine zone. *New phytologist* 11:37–50.

Kalinin V (1965) Functionals related to the poisson distribution and statistical structure of a text. *Articles on Mathematical Statistics and the Theory of Probability* pp. 202–220.

Kroes I, Lepp PW, Relman DA (1999) Bacterial diversity within the human subgingival crevice. *Proceedings of the National Academy of Sciences* 96:14547–14552.

Lu S, Zong C, Fan W, Yang M, Li J, Chapman AR, Zhu P, Hu X, Xu L, Yan L et al. (2012) Probing meiotic recombination and aneuploidy of single sperm cells by whole-genome sequencing. *Science* 338:1627–1630.

Mora C, Tittensor DP, Adl S, Simpson AG, Worm B (2011) How many species are there on earth and in the ocean? *PLoS Biol* 9:e1001127.

Ng SB, Buckingham KJ, Lee C, Bigham AW, Tabor HK, Dent KM, Huff CD, Shannon PT, Jabs EW, Nickerson DA et al. (2010) Exome sequencing identifies the cause of a mendelian disorder. *Nature genetics* 42:30–35.

Pearman PB, Weber D (2007) Common species determine richness patterns in biodiversity indicator taxa. *Biological conservation* 138:109–119.

Preston FW (1948) The commonness, and rarity, of species. *Ecology* 29:254–283.

Sichel HS (1975) On a distribution law for word frequencies. *Journal of the American Statistical Association* 70:542–547.

Stam A (1987) Statistical problems in ancient numismatics. *Statistica Neerlandica* 41:151–174.

Van Der Heijden PG, Cruyff M, Van Houwelingen HC (2003) Estimating the size of a criminal population from police records using the truncated poisson regression model. *Statistica Neerlandica* 57:289–304.

Wang J, Fan HC, Behr B, Quake SR (2012) Genome-wide single-cell analysis of recombination activity and de novo mutation rates in human sperm. *Cell* 150:402–412.

Zafarani R, Liu H (2009) Social computing data repository at ASU.

Zipf GK (1935) *The psycho-biology of language* Houghton, Mifflin.

Zong C, Lu S, Chapman AR, Xie XS (2012) Genome-wide detection of single-nucleotide and copy-number variations of a single human cell. *Science* 338:1622–1626.

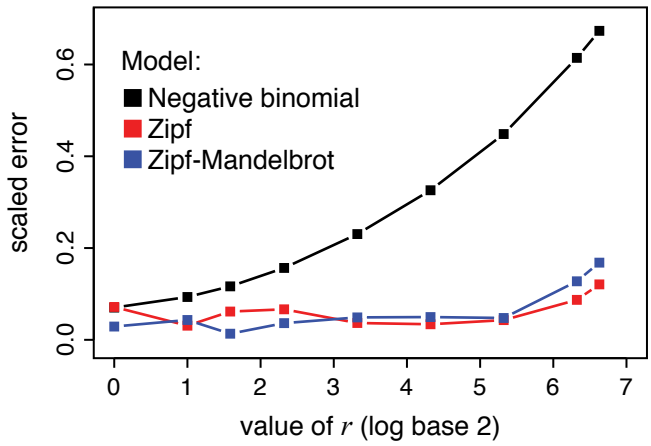

Figure 1: The scaled error as a function of r. The scaled error between the expected curve and the estimated curve as a function of r. The x-axis is in log scale (base 2) and the y-axis gives scaled error based on L_2 norm.

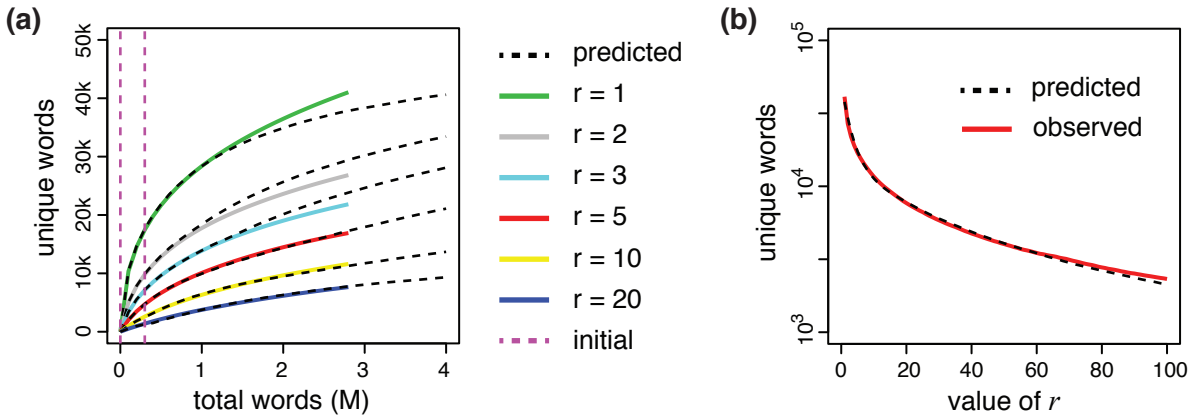

Figure 2: Words represented at least r times in samples of Dickens' writing. (**a**) Words represented at least r times as a function of total words observed. Expected values were obtained by subsampling the full data set without replacement. Estimates are based on an initial sample of 300k words, as indicated by vertical dashed lines. (**b**) Words represented at least r times along with estimated values, for $r = 1, 2, \ldots$ (y-axis is in log scale).

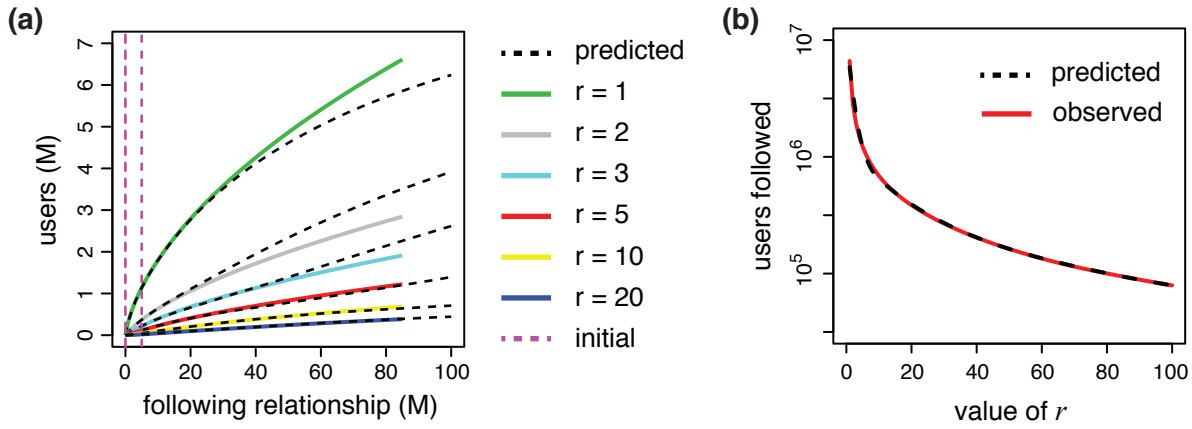

Figure 3: Twitter users with at least r followers in a sample of following relations. (**a**) Users with at least r followers as a function of total following relations. Expected values were obtained by subsampling the full data set without replacement. Estimates are based on an initial sample of 5M relations, as indicated by vertical dashed lines. (**b**) Users with at least r followers along with estimated values, for $r = 1, 2, \ldots$ (y-axis is in log scale).

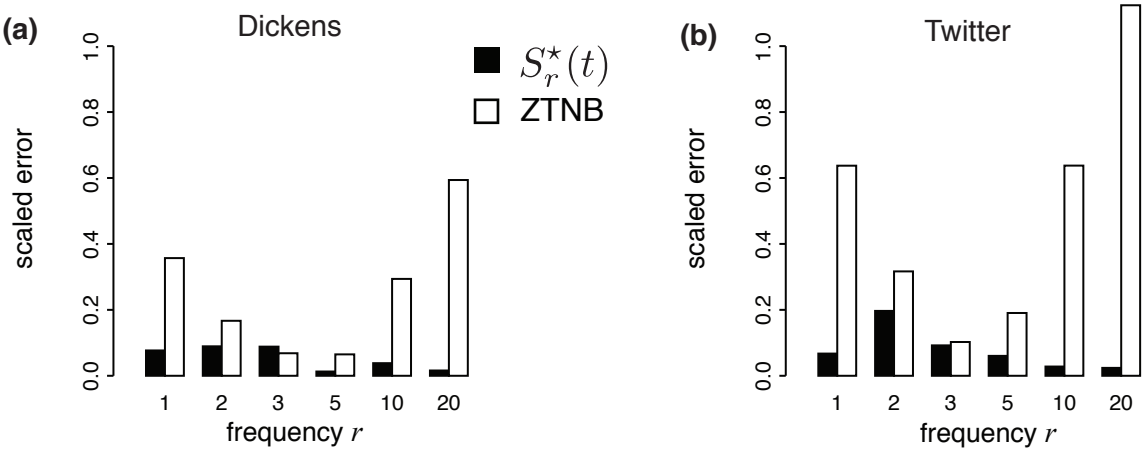

Figure 4: Scaled error of estimates for the Dickens and Twitter data sets. The scaled error is defined as the absolute of the subtraction of true value and its estimate divided by the true value. $S_r^\star(t)$, is our estimator. ZTNB, the zero-truncated negative binomial. (**a**) Scaled error of estimates of the number of words represented at least r times in the full Dickens data set based on the initial sample of size 300k. (**b**) Scaled error of estimates of the number of users with at least r followers in the full Twitter data set based on the initial sample of size 5M.

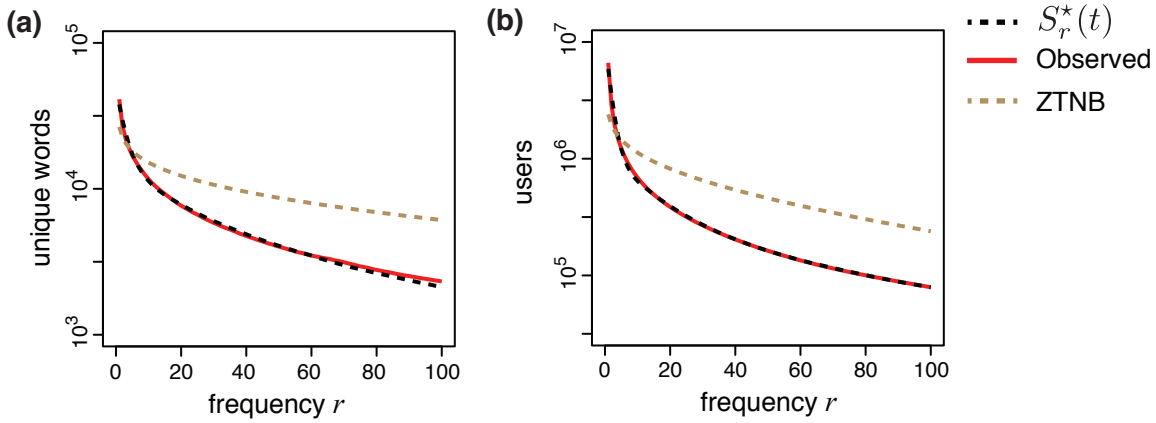

Figure 5: Accuracy of estimates as a function of r (**a**) The number of unique words represented r or more times in the Dickens data set. Estimates were made using an initial sample of size 300k. (**b**) The number of users with at least r followers in the full Twitter data set. Estimates were made using an initial sample of size 5M. Estimates using $S_r^\star(t)$ are shown as black dashed lines, estimates from ZTNB are shown as grey dashed lines, and actual expected values are shown as solid red lines.

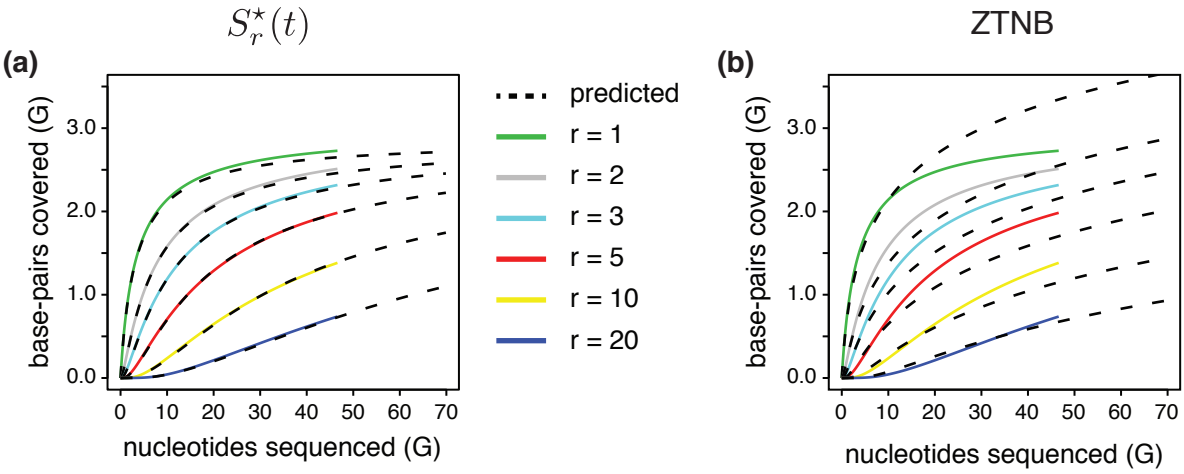

Figure 6: Base-pairs covered at least r times in a DNA sequencing data set. Expected values were obtained by subsampling the full data set without replacement. Estimates are based on an initial sample of 5M reads (500M nucleotides). Initial sample is not indicated as the small size would not be visible. Estimates made using (**a**) the estimator $S_r^\star(t)$ and (**b**) zero-truncated negative binomial (ZTNB).

	$t=2$			$t=3$			$t=5$			$t=10$			$t=20$		
r	true	mean	SE	true	mean	SE	true	mean	SE	true	mean	SE	true	mean	SE
2	62604	49805	290	68913	57732	449	75562	66503	599	82524	75797	801	87572	82156	1526
3	54249	39612	165	61707	48429	333	69726	58684	604	78254	70201	882	84500	78497	1079
5	42411	26592	180	51169	35782	207	60965	47250	376	71717	61286	794	79750	72313	1170
10	25242	11091	169	34741	18943	156	46473	30404	306	60396	46317	414	71331	60709	902
20	10339	2182	311	18119	6011	307	29817	14447	147	46041	29791	375	60107	45833	434
40	2087	112	184	5785	692	254	13997	3640	352	29501	14063	180	45824	29485	417
80	107	9	101	722	23	153	3657	276	216	13801	3477	367	29343	13869	207
99	27	7	76	279	12	126	2001	94	186	9884	1825	335	24236	9836	297

Table 1: Estimates and standard errors for the number of species represented at least r times. The actual number and the estimated number of species by $S_r^\star(t)$ based on the initial samples for the negative binomial model for $r < 100$. mean is the average estimates. SE is the standard error.

	$t=2$			$t=3$			$t=5$			$t=10$			$t=20$		
r	true	mean	SE	true	mean	SE	true	mean	SE	true	mean	SE	true	mean	SE
2	32958	32387	278	45110	44012	502	63204	60660	774	86462	82593	913	98167	99514	1869
3	19551	19366	149	28644	28522	302	44808	43994	709	72926	68237	1137	94503	89536	1260
5	9988	9731	120	14979	14697	217	24881	25176	399	48112	47618	1058	80457	73016	1392
10	4434	4406	103	6656	6572	124	11101	10621	284	22211	22519	632	44298	45457	1213
20	2095	1786	345	3147	3037	185	5253	5282	164	10516	9979	428	21042	21145	850
40	1015	540	392	1528	1101	424	2554	2347	276	5118	5172	239	10246	9703	556
80	496	142	253	749	305	333	1255	784	428	2521	2311	293	5053	5115	327
99	398	97	219	602	203	287	1010	529	398	2030	1706	376	4071	4088	266

Table 2: Estimates and standard errors of the number of species represented at least r times The expected and estimated number of species by $S_r^\star(t)$ based on the initial samples for the Zipf model for $r < 100$. mean is the average estimates. SE is the standard error.

	$t=2$			$t=3$			$t=5$			$t=10$			$t=20$		
r	true	mean	SE	true	mean	SE	true	mean	SE	true	mean	SE	true	mean	SE
2	14584	14276	206	19700	19448	571	28469	28370	1500	45344	44533	3030	67105	62883	3898
3	9129	8822	153	12469	12079	234	18436	18269	546	31069	31651	2392	50723	49731	4323
5	5478	5366	81	7490	7221	170	11105	10548	480	18939	18862	774	32235	33697	3163
10	2964	2849	84	4056	4011	125	6018	5920	189	10271	9598	690	17520	17124	822
20	1668	1346	391	2286	2089	246	3396	3332	105	5802	5761	300	9904	9204	798
40	953	515	364	1309	914	425	1949	1689	344	3337	3274	137	5701	5688	395
80	546	152	224	753	316	301	1125	701	409	1932	1670	355	3308	3245	176
99	459	101	192	635	216	258	950	510	367	1634	1303	415	2800	2692	176

Table 3: Estimates and standard errors of the number of species represented at least r times The expected and estimated number of species by $S_r^\star(t)$ based on the initial samples for the Zipf-Mandelbrot model for $r < 100$. mean is the average estimates. SE is the standard error.

		$5\times$			$9\times$	
r	observed	predicted	error	observed	predicted	error
1	32926	32283	0.02	41116	37889	0.08
2	21063	22521	0.07	26896	29505	0.10
3	16781	17199	0.02	21894	23899	0.09
5	12532	12345	0.01	16953	17345	0.02
10	8170	8205	0.00	11650	11321	0.03
20	5048	5181	0.03	7673	7794	0.02

Table 4: Accuracy in predicting Dickens' frequency of word use. Observed values and estimates from $S_r^\star(t)$ for numbers of words used at least r times in a text along with scaled error. Estimates are shown for extrapolating to a text that is $5\times$ and $9\times$ larger than the initial 300k word sample.

		$5\times$			$10\times$			$17\times$	
r	observed	predicted	error	observed	predicted	error	observed	predicted	error
1	3205418	3164608	0.01	4861765	4611910	0.05	6626985	5864172	0.12
2	1240448	1320501	0.06	1998589	2339063	0.17	2850180	3513562	0.23
3	804297	785919	0.02	1325177	1389673	0.05	1917058	2270164	0.18
5	490018	481062	0.02	831020	779502	0.06	1224301	1205101	0.02
10	251777	254193	0.01	454867	454825	0.00	688106	643717	0.06
20	122336	122847	0.00	241468	243418	0.01	386003	392038	0.02

Table 5: Accuracy in predicting number of Twitter users with multiple followers. Observed values and estimates from $S_r^\star(t)$ for numbers users with r or more followers along with scaled error. Predictions are shown for extrapolating to a social network that is $5\times$, $10\times$ and $17\times$ larger than the initial 5M following relationship sample.

arXiv:1607.02804v2 [stat.ME] 19 Aug 2016

Supplementary Information

Chao Deng Andrew D Smith

August 23, 2016

1 The proof of Proposition 1

Proposition 1. *Let $S_r(t)$ denote the number of species observed at least r times in a sample of size tN. For any positive integer r,*

$$\mathrm{E}(S_r(t)) = \frac{(-1)^{r-1}t^r}{(r-1)!}\frac{d^{r-1}}{dt^{r-1}}\left(\frac{\mathrm{E}(S_1(t))}{t}\right).$$

Proof. The proposition is trivially satisfied for $r = 1$, so we assume $r \geq 2$. By Equation (1) and linearity of expectation,

$$\mathrm{E}(S_r(t)) = \mathrm{E}(S_1(t)) - \sum_{j=1}^{r-1}\mathrm{E}(n_j(t)).$$

Equation (4) shows how to eliminate the $\mathrm{E}(n_j(t))$ and write $\mathrm{E}(S_r(t))$ in terms of derivatives of $\mathrm{E}(S_1(t))$:

$$\begin{aligned}\mathrm{E}(S_r(t)) &= \mathrm{E}(S_1(t)) - \sum_{j=1}^{r-1}\frac{(-1)^{j-1}t^j}{j!}\frac{d^j}{dt^j}\mathrm{E}(S_1(t)) \\ &= \sum_{j=0}^{r-1}\frac{(-1)^j t^j}{j!}\frac{d^j}{dt^j}\mathrm{E}(S_1(t)). \qquad \text{(S1)}\end{aligned}$$

Extracting a factor of t^r from the summands results in

$$\mathrm{E}(S_r(t)) = t^r\sum_{j=0}^{r-1}\frac{(-1)^j}{j!}\frac{1}{t^{r-j}}\frac{d^j}{dt^j}\mathrm{E}(S_1(t)).$$

Note that $1/t^{r-j}$ can be expanded as

$$\frac{1}{t^{r-j}} = \frac{(-1)^{r-1-j}}{(r-1-j)!}\frac{d^{r-1-j}}{dt^{r-1-j}}\left(\frac{1}{t}\right),$$

which suggests a substitution in Equation (S1). Eliminating the factor $1/t^{r-j}$ and introducing the $(r-1-$

$j)^{\text{th}}$ derivative of $1/t$, we arrive at:

$$\begin{aligned} t^r \sum_{j=0}^{r-1} \frac{(-1)^j}{j!} \frac{1}{t^{r-j}} \frac{d^j}{dt^j} \mathrm{E}(S_1(t)) &= t^r \sum_{j=0}^{r-1} \frac{(-1)^j}{j!} \frac{(-1)^{r-1-j}}{(r-1-j)!} \frac{d^j}{dt^j} \mathrm{E}(S_1(t)) \frac{d^{r-1-j}}{dt^{r-1-j}} \left(\frac{1}{t}\right) \\ &= t^r \sum_{j=0}^{r-1} \frac{(-1)^{r-1}}{(r-1)!} \binom{r-1}{j} \frac{d^j}{dt^j} \mathrm{E}(S_1(t)) \frac{d^{r-1-j}}{dt^{r-1-j}} \left(\frac{1}{t}\right) \\ &= \frac{(-1)^{r-1} t^r}{(r-1)!} \sum_{j=0}^{r-1} \binom{r-1}{j} \frac{d^j}{dt^j} \mathrm{E}(S_1(t)) \frac{d^{r-1-j}}{dt^{r-1-j}} \left(\frac{1}{t}\right). \end{aligned}$$

Finally, by noticing that the above summation has the form of the general Leibniz rule,

$$\begin{aligned} \mathrm{E}(S_r(t)) &= \frac{(-1)^{r-1} t^r}{(r-1)!} \sum_{j=0}^{r-1} \binom{r-1}{j} \frac{d^j}{dt^j} \mathrm{E}(S_1(t)) \frac{d^{r-1-j}}{dt^{r-1-j}} \left(\frac{1}{t}\right) \\ &= \frac{(-1)^{r-1} t^r}{(r-1)!} \frac{d^{r-1}}{dt^{r-1}} \left(\frac{\mathrm{E}(S_1(t))}{t}\right). \end{aligned}$$

□

2 Extending the domain of estimator

In the main text our estimator $S_r^\star(t)$ for the expected number of species represented at least r times in a sample of size tN was valid for $t > 1$. It is desirable to extend the domain to any sample size. When $t \leq 1$, the tN individuals may be considered to have already been observed in the initial sample. For $t \leq 1$ the expected number of unique species can then be estimated as (Heck Jr. et al., 1975, Equation 1):

$$\hat{\mathrm{E}}(S_1(t)) = \sum_j n_j - \binom{N}{tN}^{-1} \sum_j n_j \binom{N-j}{tN}. \tag{S2}$$

The generalized binomial coefficients for non-integer tN are

$$\binom{x}{y} = \begin{cases} \frac{\Gamma(x+1)}{\Gamma(y+1)\Gamma(x-y+1)}, & \text{if } x, y \geq 0, \text{ and } x - y \geq 0, \\ 0, & \text{otherwise.} \end{cases}$$

The estimator in Equation (S2) is continuous over the unit interval. We generalize the above estimation for the expected number of species observed at least r times:

$$\hat{\mathrm{E}}(S_r(t)) = \sum_j n_j - \binom{N}{tN}^{-1} \sum_j n_j \left(\sum_{i=0}^{r-1} \binom{N-j}{tN-i} \binom{j}{i} \right). \tag{S3}$$

In particular, $\hat{\mathrm{E}}(S_r(0)) = 0$ and $\hat{\mathrm{E}}(S_r(1)) = \sum_{j \geq r} n_j$. Now we define the extended estimator:

$$S_r^\star(t) = \begin{cases} \sum_j n_j - \binom{N}{tN}^{-1} \sum_j n_j \left(\sum_{i=0}^{r-1} \binom{N-j}{tN-i} \binom{j}{i} \right), & \text{if } 0 \leq t < 1, \\ \sum_{i=1}^{M} c_i \left(\frac{t}{t - x_i} \right)^r, & \text{if } t \geq 1. \end{cases} \tag{S4}$$

Because of the relationship between a Padé approximant and the corresponding power series (in this case, Equation (6)), the extended estimator is continuous at $t = 1$ for any $r \leq 2M$. The estimator (S4) therefore is well-defined and continuous in $(0, \infty)$ for $r \leq 2M$.

3 Implementation

Assume for the available initial sample the number of species represented exactly j times is denoted n_j for $j = 1, 2, \ldots$ First we construct the power series given by (6) based on these counts. This will be a polynomial because $n_j = 0$ for any sufficiently large j, and necessarily for $j > N$. Assume all coefficients of the polynomial are nonzero up to degree D. We compute the Padé approximant to this polynomial for each odd degree $d \leq D$. Among all these Padé approximants that are increasing and satisfy Proposition 2, we choose the approximant with the maximum d. If none of Padé approximants satisfy this requirement, we say the construction fails. The degree must be odd because the form of Padé approximants are restricted by Equation (7). Next we calculate all roots of x_i for the denominator $B(t)$ of the Padé approxmant. Finally, we compute the coefficients c_i of a partial fraction decomposition to the obtained Padé approximant. With the roots x_i and coefficients c_i we construct the estimator in Equation (10).

For a given polynomial, the Padé approximants is implemented by calling core functions in Preseq (Daley & Smith, 2013), which can produce a truncated continued fraction approximation. We convert that truncated continued fraction into a rational function of the form $A(t)/B(t)$ using the R package polynom (Venables et al., 2014).

After the rational function is constructed, we simplify it by canceling any factors common to both the numerator $A(t)$ and denominator $B(t)$. We use the R package *polynom* to find all the roots of $A(t)$ and $B(t)$. By decomposition of polynomial, the rational function takes the form

$$C\frac{(t-y_1)(t-y_2)\cdots(t-y_{M-1})}{(t-z_1)(t-z_2)\cdots(t-z_M)},$$

where C, y_i, and z_i are complex numbers. If y_i and z_j are approximately equal, we cancel out the factor $(t - y_i)$ from the numerator and $(t - z_j)$ from the denominator. For reasons of numerical convenience, we say two roots x_i and z_j are equal if their absolute difference is less than a threshold (10^{-3} by default). The simplified rational function is then passed to the next step (partial fractions) and the obtained estimator is evaluated using the criteria outlined in the main text.

We repeat the above procedure by bootstrapping until a specified number (100 by default) of estimators are generated. We then use the median value from all these estimators as our default final estimation generated from out implementation. A confidence interval is also calculated using the log-normal approach as suggested by Chao (1987). In each bootstrap iteration, there is a chance the algorithm will be unable to construct an estimator, in particular if all Padé approximants fail. If the number of failing iterations exceeds a threshold (the specified number divided by 0.4) before required enough estimators are constructed, we stop the algorithm and return null.

4 Estimating species represented exactly r times

The estimator $S_r^\star(t)$ can easily be used to estimate the expected number $\mathrm{E}(n_r(t))$ of species represented *exactly* r times in a sample of size tN. Having accurate estimates for the exact counts essentially amounts to predicting the entire histogram of species frequency counts. By definition,

$$n_r(t) = S_r(t) - S_{r+1}(t).$$

So to estimate $\mathrm{E}(n_r(t))$ we first estimate $\mathrm{E}(S_r(t))$ and $\mathrm{E}(S_{r+1}(t))$ according to $S_r^\star(t)$ and then take their difference.

5 Simulation

We compare the estimator $S_r^\star(t)$ for $r = 1$ with two popular parametric approaches, zero-truncated negative binomial and zero-truncated Poisson respectively, and a nonparametric approach suggested by Chao & Shen (2004). The simulation scheme is the same as it was in Chao & Shen (2004). Model 3–5 have been described in the manuscript. Model 1–2 are used to simulate homogeneous population, which are rarely adopted in practice. They are just for comparison purposes. The procedure has been described in the manuscript.

Model 1: (Homogenous model). $\lambda_i = 1.8$ for $i = 1, 2, \ldots, 100000$ ($CV = 0$, $\mathrm{E}(n) = 180,000$).

Model 2: (Two-class model). $\lambda_i = 2$ for $i = 1, 2, \ldots, 90000$, and $\lambda_i = 0.2$ for $i = 90001, 90002, \ldots, 100000$ ($CV = 0.30$, $\mathrm{E}(n) = 180,000$).

Model 3: (Negative binomial model). λ_i are generated from a gamma density

$$f(x) = \frac{1}{\Gamma(\alpha)\beta^\alpha} x^{\alpha-1} e^{-\frac{x}{\beta}}$$

with $\alpha = 0.5$, $\beta = 3.6$ ($CV = 1.41$, $\mathrm{E}(n) = 180,000$).

Model 4: (Zipf model). $\lambda_i = 20000/(i + 10)$ for $i = 1, 2, \ldots, 100000$ ($CV = 10.60$, $\mathrm{E}(n) = 183226$).

Model 5: (Zipf-Mandelbrot model). $\lambda_i = 150000/(y + 20)^{1.31}$ for $i = 1, 2, \ldots, 100000$ ($CV = 17.94$, $\mathrm{E}(n) = 186225$).

For estimating the species accumulation curve (the case $r = 1$), the scaled error decreased as the CV increased for the estimator $S_1^\star(t)$ (Figure S1). The estimator $S_1^\star(t)$ performed best when the population was highly heterogeneous, which are Model 4 and 5. The scaled error between the estimated curve and the expected curve is 0.0297 on extremely heterogeneous population (Model 5), while the second best has scaled error 0.3016. The other methods tended to saturate early (Table S5). Even for Model 3 based on a negative binomial distribution, the estimator $S_1^\star(t)$ had a comparable performance with the zero-truncated negative binomial method. However, the estimator $S_1^\star(t)$ did not work well comparing with other methods, when the population was homogeneous ($CV < 0.5$). For example on Model 2, the estimator $S_1^\star(t)$ overestimated the population size by 18,263 at $t = 10$ and overestimated the population size by 42,536 at $t = 20$, while other methods perfectly estimated the species richness (Table S2). One possible reason is that the estimator $S_1^\star(t)$ is nonparametric that has no assumption on the underlying population. The estimation favors parametric models when the population is close to their assumptions.

The standard error of the estimator $S_1^\star(t)$ is large comparing with other methods (Table S1 – S5), and it increased as a function of t. This is a bias-variance tradeoff. The estimator $S_1^\star(t)$ tends to incorporate as many high frequent counts n_j ($j \geq 4$) as possible to improve the accuracy of the prediction, which in turn brings more variance. Note that the numerator can be decomposed into the square of bias plus the square of standard error. The scaled error actually measures both the bias and the variance. From Figure S1, the estimator $S_1^\star(t)$ has best performance on highly heterogeneous population when considering both the bias and the variance.

6 DNA sequencing analysis

Four publicly available single-cell DNA sequencing datasets were downloaded from National Center for Biotechnology Information (NCBI). The Sequencing Read Archive accession numbers are SRX202787, SRX205367, SRX204160 and SRX151616. For each dataset, we subsampled 5M single-end reads as an initial sample. The sample was mapped to human genome reference GRCh38/hg38 by BWA using default

parameters (Li & Durbin, 2009). We used Picard tools (Broad Institute, 2013) to sort mapped reads and Samtools (Li et al., 2009) to obtain the number of nucleotides covered by j sequenced reads, for $j = 1, 2, \ldots$ These formed the counts n_j used to make predictions.

7 Confidence interval

We bootstrap the initial sample to generate estimators and use the log-normal approach to construct the confidence interval for $S_r^\star(t)$ Chao (1987). In general, the estimated variance of $S_r^\star(t)$ increases as the relative sample size t increases (Figure S3). To provide reasonable estimates, we restrict our estimates within 20 times the size of the initial sample for general purpose. For applications like high-throughput sequencing, we usually observed accurate estimates for extrapolating over 20 times.

We often observed a narrower region in our estimated confidence as t increases for large r. Unlike the special case $r = 1$, $\mathrm{E}(S_r(t))$ is not concave for $r > 1$. The curves based on bootstrap tend to cross each other around the inflection point, where the second derivative is 0. This causes the confidence band to become more narrow in this region, an observation that can be explained, but one to which users are likely unaccustomed. To ease visualization for users, our implementation identifies such narrowing of intervals and interpolates through this region based on wider points to the left and right. This is always conservative and maintains the stated confidence for the interval that is displayed.

		$S_r^\star(t)$		Chao & Shen		ZTNB		ZTP	
t	TRUE	AE	SE	AE	SE	AE	SE	AE	SE
2	97268	97371	260	97294	165	97321	158	97273	152
3	99548	100213	829	99587	190	99631	185	99555	168
5	99988	103905	3837	100031	198	100085	195	99994	173
10	1e+05	117687	18080	100044	198	100098	196	100006	173
20	1e+05	141716	51251	100044	198	100098	196	100006	173

Table S1: Accuracy for predicting the number of species The expected and estimated number of species based on the initial samples in model 1. AE is the average estimates. SE is the standard error. We compare the performance of four methods. $S_r^\star(t)$, Chao & Shen, ZTNB, ZTP represent our estimator, the estimator suggested by (Chao & Shen, 2004), zero-truncated negative binomial and zero-truncated Poisson.

		$S_r^\star(t)$		Chao & Shen		ZTNB		ZTP	
t	TRUE	AE	SE	AE	SE	AE	SE	AE	SE
2	91648	91597	251	91133	161	91055	162	90711	141
3	94289	94612	847	92911	188	92872	198	92338	154
5	96317	98652	3759	93220	196	93219	213	92604	157
10	98647	109540	15570	93227	196	93230	215	92610	157
20	99817	125651	39900	93227	196	93230	215	92610	157

Table S2: Accuracy for predicting the number of species The expected and estimated number of species based on the initial samples in model 2. AE is the average estimates. SE is the standard error. We compare the performance of four methods. $S_r^\star(t)$, Chao & Shen, ZTNB, ZTP represent our estimator, the estimator suggested by Chao & Shen (2004), zero-truncated negative binomial and zero-truncated Poisson.

		$S_r^\star(t)$		Chao & Shen		ZTNB		ZTP	
t	TRUE	AE	SE	AE	SE	AE	SE	AE	SE
t=2	74518	65062	233	64141	199	65080	200	55465	164
t=3	78965	70813	421	67493	242	70892	258	55547	166
t=5	83560	76774	906	68861	270	77063	347	55550	166
t=10	88296	82674	2068	69007	275	83568	484	55550	166
t=20	91695	86513	3799	69008	275	88307	622	55550	166

Table S3: Accuracy for predicting the number of species The expected and estimated number of species based on the initial samples in model 3. AE is the average estimates. SE is the standard error. We compare the performance of four methods. $S_r^\star(t)$, Chao & Shen, ZTNB, ZTP represent our estimator, the estimator suggested by (Chao & Shen, 2004), zero-truncated negative binomial and zero-truncated Poisson.

		$S_r^\star(t)$		Chao & Shen		ZTNB		ZTP	
t	TRUE	AE	SE	AE	SE	AE	SE	AE	SE
t=2	61056	60586	242	61009	223	53567	170	43185	143
t=3	72376	72058	451	72168	330	60348	200	43196	143
t=5	85147	85745	1028	82671	494	69007	239	43196	143
t=10	96245	100856	2489	87912	638	80879	293	43196	143
t=20	99680	111121	4649	88319	658	92822	348	43196	143

Table S4: Accuracy for predicting the number of species The expected and estimated number of species based on the initial samples in model 4. AE is the average estimates. SE is the standard error. We compare the performance of four methods. $S_r^\star(t)$, Chao & Shen, ZTNB, ZTP represent our estimator, the estimator suggested by (Boneh et al., 1998), (Chao & Shen, 2004), zero-truncated negative binomial and zero-truncated Poisson.

		$S_r^\star(t)$		Chao & Shen		ZTNB		ZTP	
t	TRUE	AE	SE	AE	SE	AE	SE	AE	SE
2	33098	32469	193	32402	182	26248	124	22070	102
3	41222	40425	372	39549	270	28848	140	22070	102
5	53053	51842	1116	47420	418	32141	160	22070	102
10	70658	68485	4190	52790	591	36626	188	22070	102
20	86597	84346	11116	53564	636	41122	216	22070	102

Table S5: Accuracy for predicting the number of species The expected and estimated number of species based on the initial samples in model 5. AE is the average estimates. SE is the standard error. We compare the performance of four methods. $S_r^\star(t)$, Chao & Shen, ZTNB, ZTP represent our estimator, the estimator suggested by Chao & Shen (2004), zero-truncated negative binomial and zero-truncated Poisson.

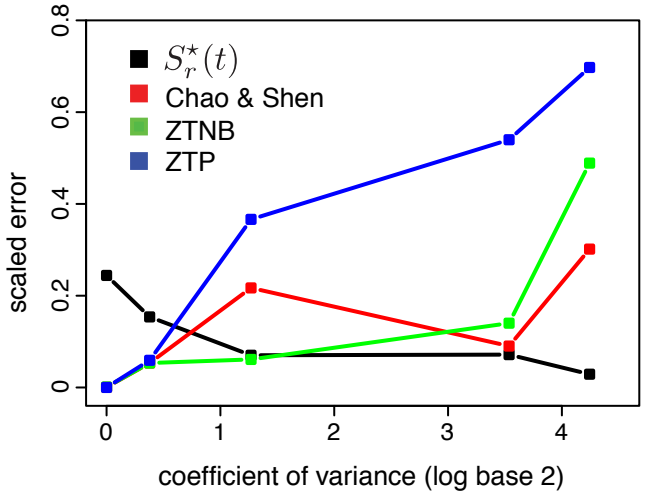

Figure S1: The scaled error as a function of the coefficient of variance (CV). The scaled error between the expected curve and the estimated curve as a function of the CV for the relative species abundance. The x-axis is the logarithm of the CV plus 1 (log (1 + CV)). The y-axis is the scaled error based on L_2 norm. $S_r^\star(t)$, Chao & Shen, ZTNB, ZTP represent our estimator, the estimator suggested by Chao & Shen (2004), zero-truncated negative binomial and zero-truncated Poisson.

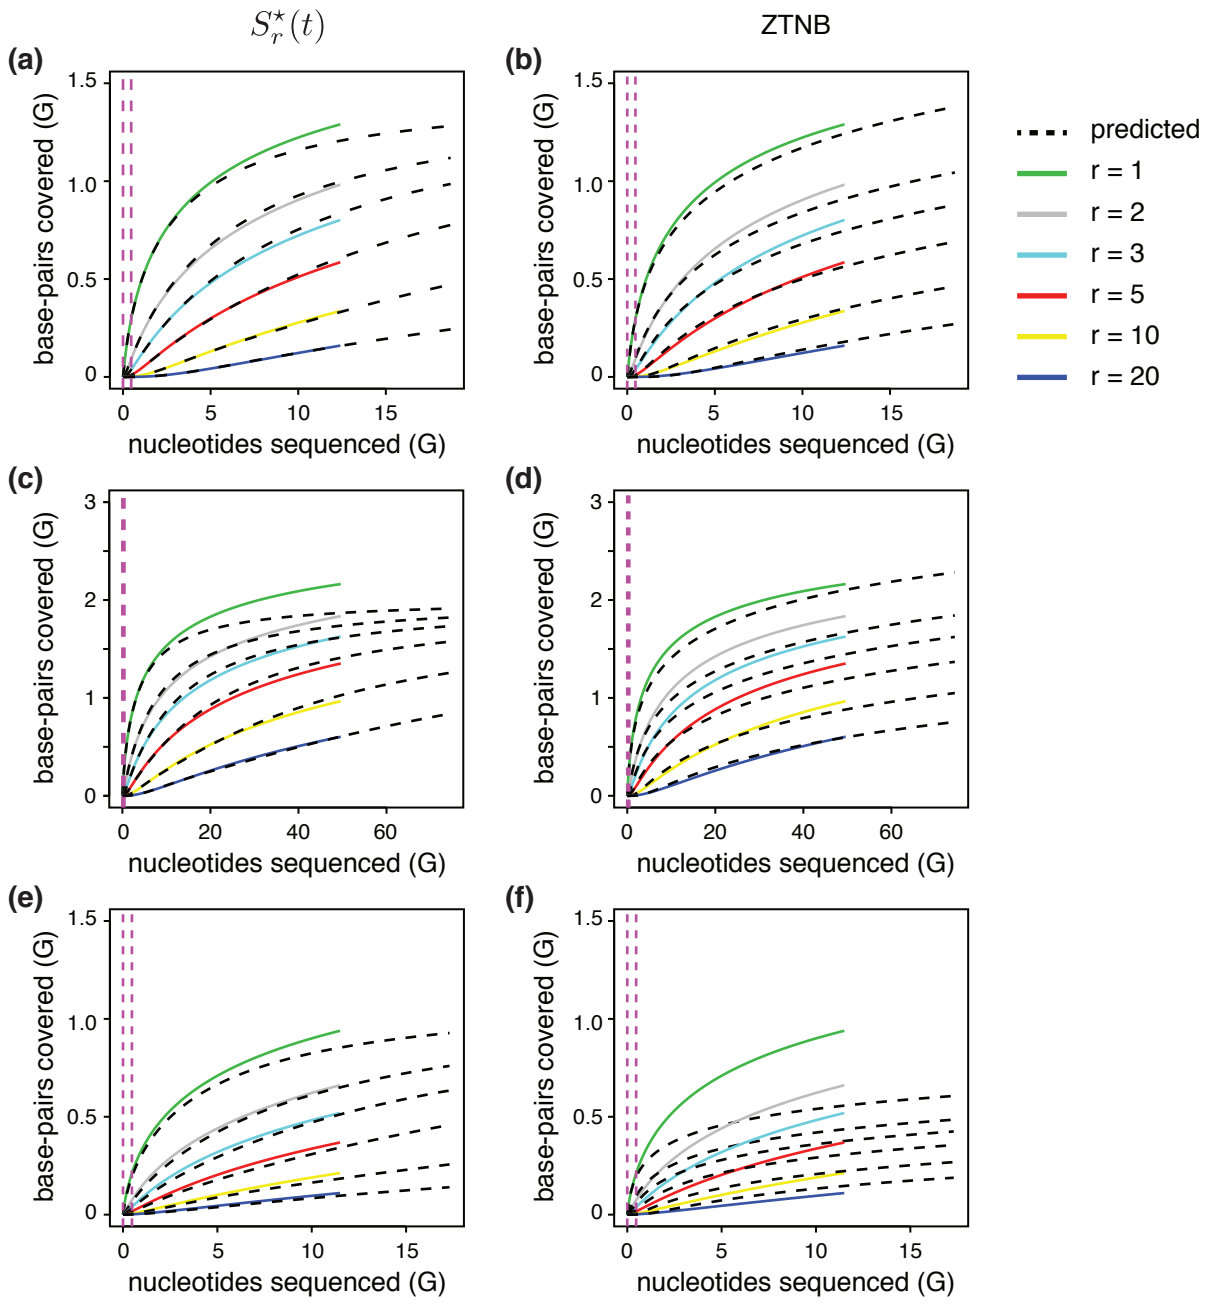

Figure S2: Accuracy of estimates on DNA sequencing data. The number of base pairs in the genome represented r or more times in the sequencing data set. Estimates were made using an initial sample of 5M reads (500M nucleotides). Estimates for data set SRX205367 made using $S_r^\star(t)$ (**a**) and ZTNB (**b**). Estimates for data set SRX204160 made using $S_r^\star(t)$ (**c**) and ZTNB (**d**). Estimates for data set SRX151616 made using $S_r^\star(t)$ (**e**) and ZTNB (**f**).

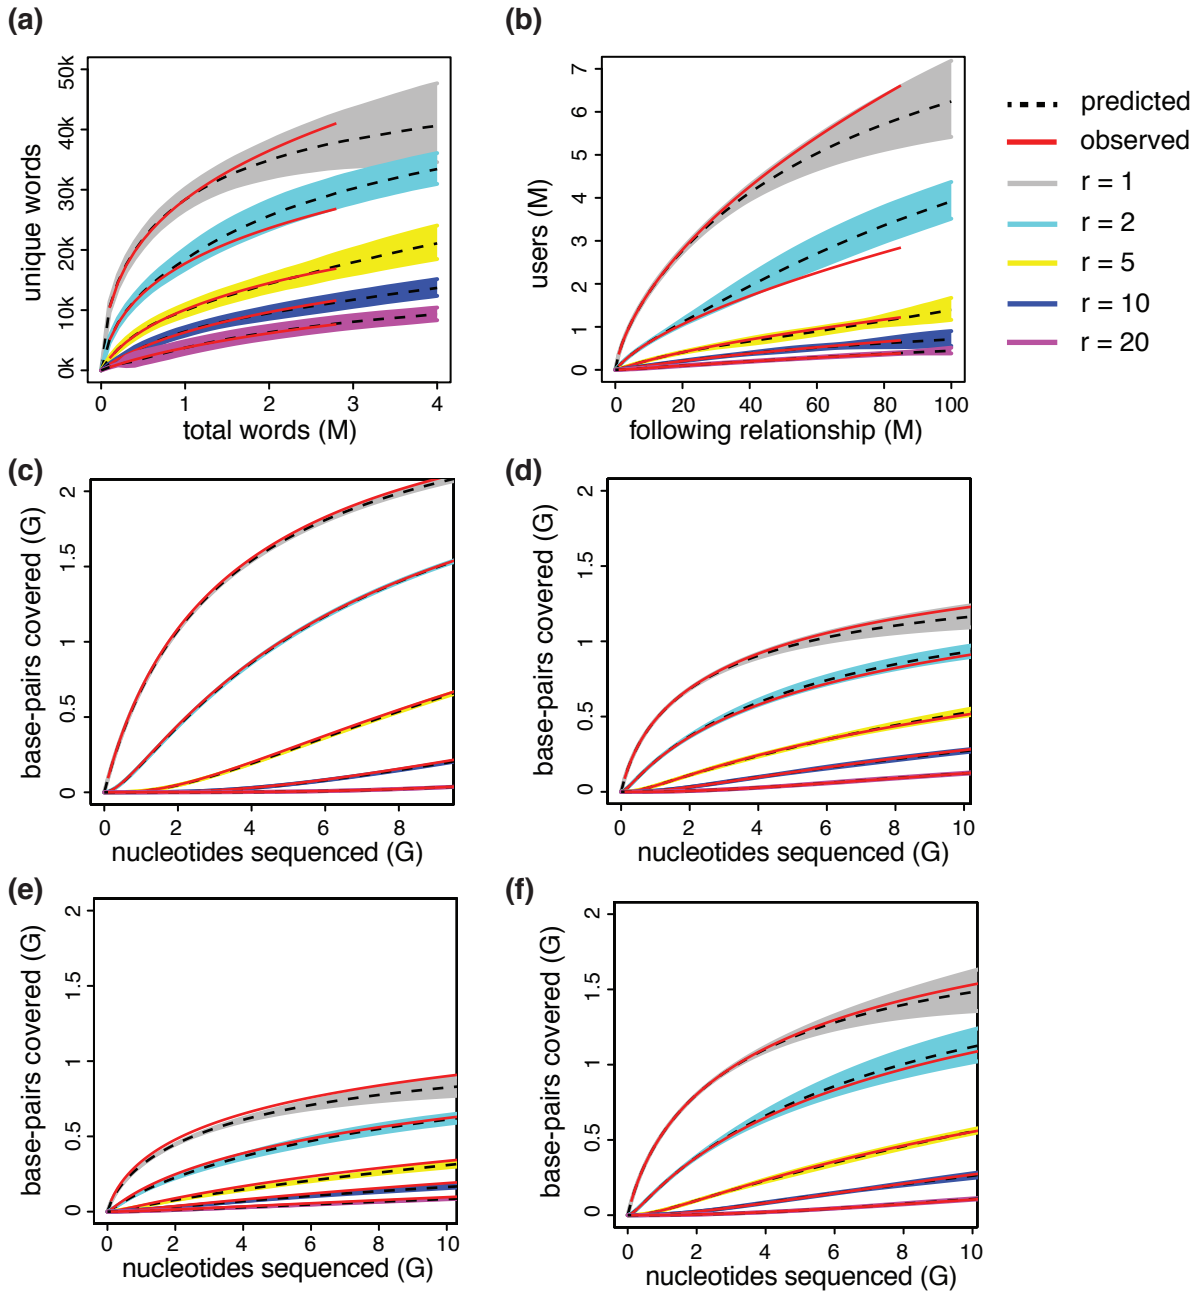

Figure S3: The estimates and their confidence bands. The estimated number of species represented r or more times and their 95% confidence interval. Estimates were made using an initial sample of **(a)** collection of Charles Dickens's works **(b)** (follower, followed) relation of Twitter users **(c)** DNA sequencing data set SRX202787 **(d)** DNA sequencing data set SRX205367 **(e)** DNA sequencing data set SRX151616 and **(f)** DNA sequencing data set SRX204160.

References

Boneh S, Boneh A, Caron RJ (1998) Estimating the prediction function and the number of unseen species in sampling with replacement. *Journal of the American Statistical Association* 93:372–379.

Broad Institute (2013) Picard Tools `http://broadinstitute.github.io/picard/`.

Chao A (1987) Estimating the population size for capture-recapture data with unequal catchability. *Biometrics* 43:783–791.

Chao A, Shen TJ (2004) Nonparametric prediction in species sampling. *Journal of agricultural, biological, and environmental statistics* 9:253–269.

Daley T, Smith AD (2013) Predicting the molecular complexity of sequencing libraries. *Nature Methods* 10:325–327.

Heck Jr. KL, van Belle G, Simberloff D (1975) Explicit calculation of the rarefaction diversity measurement and the determination of sufficient sample size. *Ecology* pp. 1459–1461.

Li H, Durbin R (2009) Fast and accurate short read alignment with Burrows–Wheeler transform. *Bioinformatics* 25:1754–1760.

Li H, Handsaker B, Wysoker A, Fennell T, Ruan J, Homer N, Marth G, Abecasis G, Durbin R et al. (2009) The sequence alignment/map format and samtools. *Bioinformatics* 25:2078–2079.

Venables B, Hornik K, Maechler M (2014) *polynom: A Collection of Functions to Implement a Class for Univariate Polynomial Manipulations* R package version 1.3-8. S original by Bill Venables, packages for R by Kurt Hornik and Martin Maechler.